AFRL-IF-RS-TR-2005-294
Final Technical Report
August 2005

GAME THEORETICAL ALGORITHMS FOR INTELLIGENT NETWORKING

University of Pennsylvania

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. N908

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2005-294 has been reviewed and is approved for publication

APPROVED: /s/

ROBERT L. KAMINSKI
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	AUGUST 2005	Final Jun 02 – Dec 03

4. TITLE AND SUBTITLE
GAME THEORETICAL ALGORITHMS FOR INTELLIGENT NETWORKING

5. FUNDING NUMBERS
C - F30602-02-2-0185
PE - 61101E
PR - N908
TA - 01
WU - 04

6. AUTHOR(S)
Michael Kerns,
Luis Ortiz and
Sham Kakade

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Pennsylvani
3451 Walnut Street
140 Franklin Building
Philadelphia Pennsylvania 19401-6270

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington Virginia 22203-1714

AFRL/IFGA
525 Brooks Road
Rome New York 13441-4505

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2005-294

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Robert L. Kaminski/IFGA/(315) 330-1865/ Robert.Kaminski@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 Words)***
This final tech report describes progress on the development of powerful new models and algorithms for very large scale games and economic problems with rich network structure. The report provides a brief synopsis of the main findings, along with pointers to relevant published papers. This effort developed a general distributed algorithm for efficiently computing approximate Nash equilibrium in large games on networks. The effort also examined the properties and computation of correlated equilibria (a powerful generalization of Nash equilibria) in large games on networks.

14. SUBJECT TERMS
Game Theory, Nash Equilibria, Security Games, Equilibrium Computation

15. NUMBER OF PAGES
13

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Table of Contents

Final Report for F30602-02-2-0185: Game-Theoretic Algorithms for Intelligent Networking

Principal Investigator: Prof. Michael Kearns
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104
mkearns@cis.upenn.edu

1 Summary

The funding was used primarily for the support of two postdoctoral associates, Drs. Luis Ortiz (www.cis.upenn.edu/~leortiz) and Sham Kakade (www.cis.upenn.edu/~skakade). Working closely with these two colleagues, Kearns made considerable progress on the development of powerful new models and algorithms for very large-scale games and economic problems with a rich network structure. A brief synopsis of the main findings, along with pointers to the relevant papers, is given below, followed by further details on the methods and results in subsequent sections.

- Kearns and Ortiz, following earlier work of Kearns and other colleagues, developed the first general distributed algorithm for efficiently computing approximate Nash equilibria in large games on networks. They established basic convergence properties of this algorithm, and demonstrated its effectiveness on a set of benchmark problems.

  Nash Propagation for Loopy Graphical Games. With L. Ortiz. Proceedings of Neural Information Processing Systems 2002. Paper is available for download at www.cis.upenn.edu/ mkearns/papers/nashprop.pdf

- Kearns and Ortiz studied algorithms for computing Nash equilibria in a recently introduced model known as Interdependent Security (IDS) games, which capture the tradeoffs between direct and indirect sources of risk in problems such as airline baggage screening. After establishing basic algorithmic properties of the model, they did an experimental study using civilian flight itinerary data.

  Algorithms for Interdependent Security Games. With L. Ortiz. Neural Information Processing 2003. Paper is available for download at www.cis.upenn.edu/ mkearns/papers/ids.pdf

- Kakade, Kearns and Ortiz (with visitor John Langford) examined the properties and computation of correlated equilibria (a powerful generalization of Nash equilibria) in large games on networks. They established a fundamental connection between the network of the game, and the structure of the minimal Markov network required to represent the correlations required at equilibrium. They also provide algorithms for correlated equilibrium computation.

Correlated Equilibria in Graphical Games. With S. Kakade, J. Langford, and L. Ortiz. ACM Conference on Electronic Commerce, 2003. Paper is available for download at www.cis.upenn.edu/ mkearns/papers/cegg.pdf

- Kakade, Kearns and Ortiz introduced a network generalization of classical models in mathematical economics, such as the Fisher and Arrow-Debreu models. They established equilibrium existence in this model, and provided among the very first algorithms for efficient approximate equilibrium computation.

  Graphical Economics. With S. Kakade and L. Ortiz. To appear, Conference on Computational Learning Theory 2004. Paper is available for download at www.cis.upenn.edu/ mkearns/papers/graphecon.pdf

2 Introduction

Recently there has been renewed interest in game theory in several research disciplines, with its uses ranging from the modeling of evolution to the design of electronic markets, network protocols, and distributed algorithms. In the AI community, game theory is emerging as the dominant formalism for studying strategic and cooperative interaction in multi-agent systems. Classical work provides rich mathematical foundations and equilibrium concepts, but relatively little in the way of computational and representational insights that would allow game theory to scale up to large, complex systems.

One promising direction for making computational progress is to introduce alternative ways of *representing* these problems, with the hope that wide classes of "natural" problems may permit special-purpose solutions. By developing new representations that permit the expression of common types of structure in games and economies, it may be possible to design algorithms that exploit this structure to yield computational as well as modeling benefits. Researchers in machine learning and artificial intelligence have proven especially adept at devising models that balance representational power with computational tractability and learnability, so it has been natural to turn to these literatures for inspiration in strategic and economic models.

Among the most natural and common kinds of structure that arise in game-theoretic and economic settings are constraints and asymmetries in the *interactions* between the parties. By this we mean, for example, that in a large-population game, not all players may directly influence the payoffs of all others. The recently introduced formalism of *graphical games* captures this notion, representing a game by an undirected graph and a corresponding set of local game matrices (Kearns et al. [2001]).

Graphical games are a representation consisting of an undirected graph and a set of local payoff matrices was proposed for multi-player games. The interpretation is that the payoff to player i is a function of the actions of only those players in the neighborhood of vertex i in the graph. Exactly as with the graphical models for probabilistic inference that inspired them (such as Bayesian and Markov networks), graphical games provide an exponentially more succinct representation in cases where the number of players is large, but the degree of the interaction graph is relatively small. Also as for probabilistic inference, the problem of computing equilibria on arbitrary graphs is intractable in general, and so it is of interest to identify both natural special topologies permitting fast Nash computations, and good heuristics for general graphs.

Kearns et al. [2001] (abbreviated KLS in the sequel) gave a provably efficient (polynomial in the model size) algorithm for computing all approximate Nash equilibria in graphical games in which the underlying graph is a tree. This algorithm can be formally viewed as the analogue of the junction tree algorithm for inference in tree-structured networks (Pearl [1988]). A related algorithm described in Littman et al. [2002] computes a single but exact Nash equilibrium, again for a tree topology.

The funding here was used to extend this work to more general settings. We now outline these settings.

2.1 Generalizing from Trees to "Loopy" Graphs

A natural question following from KLS is whether there are generalizations of the basic tree algorithm analogous to those for probabilistic inference. In probabilistic inference, there are two main approaches to generalizing the polytree algorithm. Roughly speaking, the first approach is to take an arbitrary graph and "turn it into a tree" via triangulation, and subsequently run the tree-based algorithm on the resulting *junction tree* (Lauritzen and Spiegelhalter [1988]). This approach has the merit of being guaranteed to perform inference correctly, but the drawback of requiring the computation to be done on the junction tree. On highly loopy graphs, junction tree computations may require exponential time. The other broad approach is to simply run (an appropriate generalization of) the polytree algorithm on the original loopy graph. This method garnered considerable interest when it was discovered that it sometimes performed quite well empirically, and was closely connected to the problem of decoding in Turbo Codes. Belief propagation has the merit of each iteration being quite efficient, but the drawback of having no guarantee of convergence in general — though recent theoretical work has established convergence for certain special cases (Weiss [2000]).

In recent work, Vickrey and Koller [2002] proposed a number of heuristics for equilibria computation in graphical games, including a constraint satisfaction generalization of KLS that essentially provides a junction tree approach for arbitrary graphical games. They also gave promising experimental results for this heuristic on certain loopy graphs that result in manageable junction trees.

In this work, we introduced the NashProp algorithm, a different KLS generalization which provides an approach analogous to loopy belief propagation for graphical games.

2.2 Computation in Interdependent Security Games

Inspired by events ranging from 9/11 to the collapse of the accounting firm Arthur Andersen, economists Kunreuther and Heal [2003] recently introduced an interesting game-theoretic model for problems of *interdependent security (IDS)*, in which a large number of players must make individual investment decisions related to security — whether physical, financial, medical, or some other type — but in which the ultimate safety of each participant may depend in a complex way on the actions of the entire population. A simple example is the choice of whether to install a fire sprinkler system in an individual condominium in a large building. While such a system might greatly reduce the chances of the owner's property being destroyed by a fire originating *within* their own unit, it might do little or nothing to reduce the chances of damage caused by fires originating in *other* units (since sprinklers can usually only douse small fires early). If "enough" other unit owners have not made the investment in sprinklers, it may be not cost-effective for any individual to do so.

Kunreuther and Heal [2003] observe that a great variety of natural problems share this basic interdependent structure, including investment decisions in airline baggage security (in which investments in new screening procedures may reduce the risk of directly checking suspicious cargo, but nearly all airlines accept transferred bags with no additional screening [1]); risk management in corporations (in which individual business units have an incentive to avoid high-risk or illegal activities only if enough other units are similarly well-behaved); vaccination against infectious disease (where the fraction of the population choosing vaccination determines the need for or effectiveness of vaccination); certain problems in com-

[1] El Al airlines is the exception to this.

puter network security; and many others. All these problems share the following important properties:

- There is a "bad event" (condominium fire, airline explosion, corporate bankruptcy, infection, etc.) to be avoided, and the opportunity to reduce the risk of it via some kind of investment.
- The cost-effectiveness of the security investment for the individual is a function of the investment decisions made by the others in the population.

The original work by Kunreuther and Heal [2003] proposed a parametric game-theoretic model for such problems, but left the interesting question of *computing* the equilibria of model largely untouched. In this work, we examined such computational issues.

2.3 An Alternative Equilibrium Notion: Correlated Equilibrium

Like much of the history of game theory, the study of graphical games so far has been dominated by Nash's classical notion of equilibrium, for which it always suffices to consider product distributions over the players' joint actions. Here, we examined *correlated equilibria* (CE) (Aumann [1974]) in graphical games, which allow arbitrary joint distributions.

In classical game theory, correlated equilibria offer a number of conceptual and computational advantages over Nash equilibria. One of the most interesting aspects of CE is that they broaden the set of "rational" solutions for normal form games without the need to address often difficult issues such as stability of coalitions and payoff imputations (Aumann [1987]). Other advantages include the facts that new and sometimes more "fair" payoffs can be achieved, that correlated equilibria can be computed efficiently for games in standard normal form, and that correlated equilibria are the convergence notion for several natural learning algorithms (Foster and Vohra [1999]). Furthermore, it has been argued that correlated equilibria are the natural equilibrium concept consistent with the Bayesian perspective (Aumann [1987], Foster and Vohra [1997]).

Correlated equilibria (CE) can be viewed as distributions over joint actions in which players are still taking best responses. Intuitively, in a CE, the action played by any player is a best response (in the expected payoff sense) to the conditional distribution over the other players given that action, and thus no player has a *unilateral* incentive to deviate from playing their role in the CE. Note that a CE may be an arbitrarily complex joint distribution. In contrast, a *Nash equilibrium* (Nash [1951]) is a special case of CE in which we demand that every player acts independently of all others.

The traffic signal is often cited as an informal everyday example of CE, in which a single bit of shared information allows a fair split of waiting times (Owen [1995]). In this example, no player stands to gain greater payoff by unilaterally deviating from the correlated play, for instance by "running a light".

We presented a series of fundamental results examining the representational and computational issues that arise when considering correlated equilibria in the compact language of graphical games.

2.4 Generalizing to the Economic Setting

Models for the exchange of goods and their prices in a large economy have a long and storied history within mathematical economics, dating back more than a century to the work of Walras [1874] and Fisher [1891], and continuing through the model of Wald [1936] (see also Brainard and Scarf [2000]). A pinnacle of this line of work came in 1954, when Arrow and Debreu provided extremely general conditions for the existence of an equilibrium in such models (in which markets clear, *i.e.* supply balances demand, and all individual consumers and firms optimize their utility subject to budget constraints). Like Nash's roughly contemporary proof of the existence of equilibria for normal-form games (Nash [1951]), Arrow and Debreu's result placed a rich class of economic models on solid mathematical ground.

These important results established the *existence* of various notions of equilibria. The *computation* of game-theoretic and economic equilibria has been a more slippery affair. Indeed, despite decades of effort, the computational complexity of computing a Nash equilibrium for a general-sum normal-form game remains unknown, with the best known algorithms requiring exponential time in the worst case. Even less is known regarding the computation of Arrow-Debreu equilibria. Only quite recently, a polynomial-time algorithm was discovered for the special but challenging case of linear utility functions (Devanur et al. [2002], Jain et al. [2003], Devanur and Vazirani [2003]).

In this work, we introduced a graph-theoretic generalization of classical Arrow-Debreu economics that permits the expression of common types of economic structure. We then provided algorithms which exploit this structure to yield computational benefits.

3 Methods and Results

We now discuss the methods used and results we obtained.

3.1 Computing Nash Equilibria on Loopy Graphs

In this work, we introduced the NashProp algorithm, a different KLS generalization which provides an approach analogous to loopy belief propagation for graphical games. Like belief propagation, NashProp is a local message-passing algorithm that operates directly on the original graph of the game, requiring no triangulation or moralization operations. NashProp is a two-phase algorithm. In the first phase, nodes exchange messages in the form of two-dimensional tables. The table player U sends to neighboring player V in the graph indicates the values U "believes" he can play given a setting of V and the information he has received in tables from his other neighbors, a kind of conditional Nash equilibrium. In the second phase of NashProp, the players attempt to incrementally construct an equilibrium obeying constraints imposed by the tables computed in the first phase.

We provided rather strong theory for the first phase, proving that the tables must always converge, and result in a reduced search space that can never eliminate an equilibrium. When run using a discretization scheme introduced by KLS, the first phase of NashProp will actually converge in time polynomial in the size of the game representation.

We also reported on a number of controlled experiments with NashProp on loopy graphs, including some that would be difficult via the junction tree approach due to the graph topology. The results appear to be quite encouraging, thus growing the body of heuristics available for computing equilibria in compactly represented games.

3.2 Computation in Interdependent Security Games

We used the IDS model given by Kunreuther and Heal [2003], in which IDS problems are modeled as n-player games. In an *IDS game*, each player must decide whether or not to invest in some abstract security mechanism (such as a sprinkler system) or procedure that can reduce their risk of experiencing some abstract bad event (such as their condominium burning down). There is a monetary cost of the investment to each player, and there is a cost of experiencing the bad event. Each, player has the choice to invest or not. Thus, there is a natural tension: players can either invest in security, which costs money but reduces risk, or gamble by not investing.

In our work we established the first computational theory of this broad class of IDS games, including efficient algorithms for equilibrium computation in cases where certain natural symmetries hold, as well as intractability results for certain equilibria in the general case.

As an empirical demonstration of IDS games, we constructed and conducted experiments on an IDS game for airline security that is based on real industry data. We have access to a data set consisting of 35,362 records of actual civilian commercial flight reservations, both domestic and international, made on August 26, 2002. Since these records contain complete flight itineraries, they include passenger transfers between the 122 represented commercial air carriers. As described below, we used this data set to construct an IDS game in which the players are the 122 carriers, the "bad event" corresponds to a bomb exploding in a bag being transported in a carrier's airplane, and the transfer event is the physical transfer of a bag from one carrier to another.

The resulting simulations exhibited a number of striking results and predictions, including the fact that the largest carriers tend not to invest in increased security, and the existence of a small "tipping set" of carriers whose subsidized or coerced investment causes the entire population to evolve towards investment.

3.3 Correlated Equilibria in Graphical Games

The first issue that arose in this investigation is the problem of *representing* correlated equilibria. Unlike Nash equilibria, even in very simple graphical games there may be correlated equilibria of essentially arbitrary complexity (for instance, any mixture distribution of Nash equilibria is a correlated equilibrium). Since one of our primary goals is to maintain the succinctness of graphical games, some way of addressing this distributional complexity is required. For this we turned to another graphical formalism — namely, undirected graphical models for probabilistic inference, also known as *Markov networks.*

Our main results established a natural and powerful relationship between a graphical game and a certain associated Markov network. It showed that the associated Markov network is sufficient for representing *all* correlated equilibria of the graphical game, up to expected payoff equivalence. Like the graphical game, the associated Markov network is a graph over the players. In other words, the fact that a multiplayer game can be succinctly represented by a graph implies that its entire space of correlated equilibria, up to payoff equivalence, can be represented graphically with comparable succinctness. Furthermore, while the interactions between vertices in the graphical game are entirely *strategic* and given by local payoff matrices, the interactions in the associated Markov network are entirely *probabilistic* and given by local potential functions. This basic result establishes a new sense in which graphical games are a powerful formalism, and highlights the natural relationship between computational game theory and modern probabilistic modeling.

Our second main result establishes the *computational* benefits of this relationship. The fact that correlated equilibria are characterized by a set of linear inequalities is not helpful for unrestricted multiplayer games, since in general there are an exponential number of such inequalities. Here again the graphical representations reap benefits. We show that a graphical game gives rise to a small set of linear inequalities (comparable in size to the game representation itself) with a non-empty feasible region that includes all correlated equilibria. In the case that the associated Markov network is *chordal,* which includes graphical game trees as a special case, we prove that any point in the feasible region can be efficiently mapped to a correlated equilibrium on the associated Markov network, thus yielding a polynomial time algorithm for computing correlated equilibria in a large class of graphical games. This algorithm also applies generally to non-chordal graphs, but in the worst case may require an exponential increase in the Markov network size.

3.4 Graphical Economics

In the same spirit of graphical games, in this work we introduced a new model called *graphical economics* and show that it provides representational and algorithmic benefits for Arrow-Debreu economics. Each vertex i in an undirected graph represents an individual party in a large economic system. The presence of an edge between i and j means that free trade is allowed between the two parties, while the absence of this edge means there is an embargo or other restriction on direct trade. The graph could thus represent a network of individual business people, with the edges indicating who knows whom; or the global economy, with the edges representing nation pairs with trade agreements; and many other settings. Since not all parties may directly engage in trade, the graphical economics model permits (and realizes) the emergence of *local* prices — that is, *the price of the same good may vary* across the economy. Indeed, one of our motivations in introducing the model is to capture the fact that price differences for identical goods can arise due to the network structure of economic interaction.

The graphical economics model suggests a *local* notion of clearance, directly derived from that of the Arrow-Debreu model. Rather than asking that the entire (global) market clear in each good, we can ask for the stronger "provincial" conditions that the *local* market for each good must clear. For instance, the United States is less concerned that the worldwide production of beef balances worldwide demand than it is that the production of *American* beef balances *worldwide* demand for American beef. If this latter condition holds, the American beef industry is doing a good job at matching the global demand for their product, even if other countries suffer excess supply or demand. Similarly, an individual salesperson in a large network for a product is principally concerned with managing just their own inventory (supply and demand).

In addition to the introduction of the graphical economics model for capturing structured interaction between individuals, organizations or nations, our contributions were in proving that such an equilibrium always exists and in providing computational algorithms. Our first result is a proof that under very general conditions (essentially analogous to Arrow and Debreu's original conditions), graphical equilibria always exist. This proof requires a non-trivial modification to that of Arrow and Debreu , due to the need to explicitly prove non-zero prices in all local neighborhoods. It combines an interesting concept known as a "quasi-equilibrium" due to Debreu [1962] with connectivity properties of the graph.

We also provided algorithm for computing approximate standard market equilibria in the non-graphical setting that runs in time polynomial in the number of players (fixing the number of goods) for a rather general class of non-linear utility functions. This result generalizes the algorithm of Deng et al. [2002] for linear utility functions. The approximation algorithms we derive for this result are more generally useful for computation of graphical equilibria. We note that while this result was not our principal interest, it fell naturally out of our pursuit of the next result (which uses it as a subroutine), and seems sufficiently fundamental to warrant highlighting.

Our final contribution here was an algorithm, called **ADProp** (for *Arrow-Debreu Propagation*) for computing approximate graphical equilibria. This algorithm is a message-passing algorithm working directly on the graph, in which neighboring consumers or economies exchange information about trade imbalances between them under potential equilibria prices. In the case that the graph is a tree, the running time of the algorithm is exponential in the graph degree and number of goods k, but only polynomial in the number of vertices n (consumers or economies). It thus represents dramatic savings over treating the graphical case with a non-graphical algorithm, which results in a running time exponential in n (as well as in k).

References

R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1, 1974.

R.J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55, 1987.

William C. Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper 1272*, August 2000.

Gerard Debreu. New concepts and techniques for equilibrium analysis. *International Economic Review*, 3(3):257–273, September 1962.

Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of equilibria. In *Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing*, pages 67–71. ACM Press, 2002.

Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002.

Nikhil R. Devanur and Vijay V. Vazirani. An improved approximation scheme for computing Arrow-Debreu prices for the linear case. In *Proceedings of the 23rd Conference on Foundations of Software Technology and Theoretical Computer Science*, 2003. To appear.

Irving Fisher. PhD thesis, Yale University, 1891.

D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 1997.

D. Foster and R. Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, pages 7 – 36, 1999.

Kamal Jain, Mohammad Mahdian, and Amin Saberi. Approximating market equilibria. In *Proceedings of the 6th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2003.

M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.

Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty (Special Issue on Terrorist Risks)*, 2003. In press.

S. Lauritzen and D. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *J. Royal Stat. Soc. B*, 50(2):157–224, 1988.

M. Littman, M. Kearns, and S. Singh. An efficient exact algorithm for singly connected graphical games. In *Neural Information Processing Systems*, 2002.

J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

G. Owen. *Game Theory*. Academic Press, UK, 1995.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.

D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002. To appear.

Abraham Wald. Über einige Gleichungssysteme der mathematischen Ökonomie (On some systems of equations of mathematical economics). *Zeitschrift für Nationalökonomie*, 7 (5):637–670, 1936. English translation by Otto Eckstein in *Econometrica*, Vol. 19, No. 4 (Oct., 1951), 368–403.

Léon Walras. *Éléments d'économie politique pure; ou, Théorie de la richesse sociale (Elements of Pure Economics, or the Theory of social wealth)*. Lausanne, Paris, 1874. (1899, 4th ed.; 1926, rev ed., 1954, Engl. transl.).

Y. Weiss. Correctness of local probability propagation in graphical models with loops. *Neural Computation*, 12(1):1–41, 2000.